

What Really Happens in the Hallways

A Multidisciplinary Analysis of Pedestrian Flows in a New Secondary School Building

Bram Leisink Myrthe Sas

Mondial College, Nijmegen, The Netherlands

Supervisor: Henk Kerssies

Abstract

Mondial College Leuvenbroek in Nijmegen is relocating to a new building, creating a rare opportunity to analyse pedestrian congestion before the move — now only months away — makes most design decisions irreversible. A three-phase study was conducted: a controlled behavioural experiment with $N = 208$ students across seven corridor configurations; a graph-theoretic simulation ($|V| = 388$, $|E| = 427$) driven by the current Leuvenbroek timetable translated to the new building's room layout, generating more than 52,000 weekly origin-destination pairs; Simulated Annealing optimisation of the room allocation; and an ongoing microscopic agent simulation using JuPedSim. The dominant finding is that 61.99% of all student movements pass through one central hall, with a Gini coefficient of 0.585 across edge loads — a consequence of the building's topology, not of room assignments or timetable structure. SA optimisation reduced load inequality by only 3%, confirming that scheduling cannot resolve architecturally determined congestion. The most effective interventions are adopting longer lesson blocks, ensuring timely classroom opening, and keeping every staircase accessible — illustrated by a scenario in which closing two outer staircases would push significantly more load onto an already overloaded central hub. The methodology is building-agnostic; post-occupancy validation is planned for September 2026.

Keywords: crowd dynamics, pedestrian simulation, graph theory, school building design, Simulated Annealing, behavioural modelling, pre-occupancy analysis

Contents

1	Introduction	1
2	Theoretical Framework	2
2.1	Pedestrian Modelling Approaches	2
2.2	Simulated Annealing	3
3	Methods	3
3.1	Behavioural Experiment	3
3.2	Building Graph Construction	3
3.3	Timetable-Driven Simulation	4
3.4	Room Allocation Optimisation	4
3.5	Microscopic Simulation Framework	4
4	Results	5
4.1	Behavioural Characterisation	5
4.2	Geometric and Topological Analysis	5
4.3	Room Allocation Optimisation	6
5	Conclusions	6
6	Discussion	7
6.1	Limitations	7
6.2	Practical Implementability	7
6.3	Broader Implications	8
6.4	Future Directions	8

1. Introduction

Mondial College Leuvenbroek in Nijmegen is relocating to a new building at the Energieweg 93. When the building opens, approximately 1,200 students will move between classrooms within a three-minute window at each lesson change. This combination — a well-defined population, a documented schedule,

and a building not yet occupied — creates a rare opportunity for *pre-occupancy analysis*: the systematic prediction of pedestrian flow from architectural drawings and empirical behavioural data, before the opportunity to act closes.

The current Leuvenbroek building shows what is at stake: corridor gridlock at lesson changes is a daily occurrence, contributing to student stress, physical friction, and late arrivals to class. Addressing these bottlenecks in the design phase, rather than after the building opens, is the central motivation for this study.

Pre-occupancy crowd studies are uncommon in practice and largely absent for school buildings in the scientific literature. Most crowd dynamics research has been developed for transport hubs, stadiums, and public spaces [1]. Schools differ from these environments in ways that matter for modelling.

First, movement in a school is governed by a timetable rather than by individual choice. Every lesson change generates a deterministic set of origin-destination pairs known in advance: student i must travel from room a to room b at time t . The result is strongly concentrated, predictable peak flows that differ fundamentally from the voluntary, distributed movement patterns in stations or shopping areas.

Second, the density contrast across a school day is extreme. During a lesson, corridors are nearly empty; at the moment the bell rings, all 1,200 students begin moving simultaneously. The ratio of peak to baseline density exceeds $\rho_{\text{peak}}/\rho_{\text{base}} > 10$. This periodic impulse structure violates the stationarity assumptions on which most macroscopic pedestrian flow models rest.

Third, students do not move as independent agents. They move as socially cohesive friend groups, with correlated routing decisions and coupled velocities [2]. Standard pedestrian models treat each agent as an autonomous particle; applying that assumption to a school population introduces a systematic error in the spatial distribution of corridor load.

Together, these three properties mean that models calibrated on adults in public spaces cannot be transferred to a school population without modification. They also mean that the consequences of architectural decisions — corridor widths, staircase placement, hall geometry — may be qualitatively different in a school than the same analysis would suggest for a station concourse.

The central question of this study is: to what extent can a combination of empirical behavioural characterisation, graph-theoretic building modelling, and algorithmic optimisation predict and mitigate pedestrian congestion before a school building enters operation?

The study proceeds in three phases. Phase 1 is a controlled behavioural experiment with $N = 208$ students across seven corridor configurations, producing the calibration parameters needed to adapt standard pedestrian models to this population. Phase 2 converts the architectural drawings of the new building into a weighted graph, drives it with the real school timetable to generate corridor load maps, and applies Simulated Annealing to evaluate room-allocation alternatives. Phase 3 constructs a microscopic agent-based simulation using JuPedSim, calibrated with the Phase 1 parameters, to resolve local phenomena (queue formation, doorway conflicts, density waves) invisible to the graph model; this phase is ongoing.

The scope of this paper is internal pedestrian flows during lesson changes. Evacuation scenarios, outdoor spaces, and non-timetabled movement are not considered.

2. Theoretical Framework

2.1 Pedestrian Modelling Approaches

Pedestrian flow models fall into three broad classes: macroscopic, microscopic, and graph-theoretic. The choice among them is not merely one of preference; each makes assumptions that either hold or fail for a given setting. All three classes were evaluated for this study, and the microscopic model was implemented and tested rather than rejected on theoretical grounds alone.

2.1.1 Macroscopic models

Macroscopic approaches treat a crowd as a continuous fluid described by a density field $\rho(\mathbf{x}, t)$ and velocity field $\mathbf{v}(\mathbf{x}, t)$, governed by the continuity equation $\partial_t \rho + \nabla \cdot (\rho \mathbf{v}) = 0$ [3]. The relation between density and speed is captured by the fundamental diagram. As an illustrative example, under the Greenshields approximation,

$$v(\rho) = v_0 \left(1 - \frac{\rho}{\rho_{\max}} \right), \quad (1)$$

with free-flow speed $v_0 \approx 1.3$ m/s and jam density $\rho_{\max} \approx 5$ pers/m², one can derive a theoretical capacity bound for a corridor. The maximum specific flow rate occurs at the optimal density $\rho_{\text{opt}} = \rho_{\max}/2 = 2.5$ pers/m², yielding a peak capacity flux of $J_{\max} = \frac{1}{4} v_0 \rho_{\max} \approx 1.625$ pers/(m · s). For the building's nominal corridor widths of 2 to 4 m, this implies an absolute theoretical limit of 3.25 to 6.50 pers/s per corridor cross-section. These numerical values depend heavily on the chosen fundamental diagram and are best treated as order-of-magnitude references rather than precise predictions for this specific building.

As a primary simulation model here, the macroscopic approach is unsuitable on two grounds. The corridor widths in the new building range from 2 to 4 m; at these scales the continuum assumption does not hold. Furthermore, students travel to heterogeneous destinations distributed across the building, so flow direction is not uniform across any cross-section. Macroscopic models are retained solely as a source of theoretical corridor capacity bounds.

2.1.2 Microscopic models — Social Force Model

The Social Force Model (SFM) [4] describes each pedestrian i as an agent whose motion is governed by

$$m_i \frac{d\mathbf{v}_i}{dt} = \underbrace{\frac{m_i}{\tau_i} (\mathbf{v}_i^0 - \mathbf{v}_i)}_{\text{goal-directed}} + \underbrace{\sum_{j \neq i} \mathbf{F}_{ij}^{\text{rep}}}_{\text{agent-agent}} + \underbrace{\sum_w \mathbf{F}_{iw}^{\text{wall}}}_{\text{wall repulsion}}, \quad (2)$$

where the goal-directed term accelerates the agent toward its desired velocity with relaxation time $\tau \approx 0.5$ s, and the repulsion terms model personal space via exponentially decaying forces. Without explicit programming, the SFM reproduces collective phenomena such as lane formation in counterflow and oscillations at bottlenecks.

The SFM was implemented in Python as part of this project and evaluated on small scenarios — single corridors, doorways, a corner — where it qualitatively reproduced the expected patterns. At building scale, with $N \approx 1,200$ students and $|V| = 388$ graph nodes, the model was computationally infeasible. Each time step requires evaluating $O(N^2)$ force pairs; spatial hashing reduces this to $O(N)$ in practice, but the continuous state space and the time step required for numerical stability ($\Delta t \approx 0.05$ s) result in simulation times of several hours per lesson change. Converting the BIM floor plans into a walkable 2D surface also presented a substantial technical barrier. Whole-building microscopic simulation is, however, becoming computationally feasible as the research progresses: optimised implementations of the SFM (as in JuPedSim) and improved computing resources are the basis of our ongoing Phase 3 work. The SFM is therefore retained not as a routing model for the current phase but as a qualitative interpretive framework and as the engine for Phase 3.

2.1.3 Graph-theoretic models

The building is represented as a weighted graph $G = (V, E, w)$, where vertices V correspond to rooms, corridor junctions, and staircase landings; edges E correspond to passable connections between them; and edge weights encode traversal time as a function of distance and walking speed. Given an origin-destination pair (s, t) , Dijkstra's algorithm yields the shortest temporal path in $O(|E| + |V| \log |V|)$ time [5]. Routing N students across a full lesson change costs $O(N \log |V|)$, making a complete weekly simulation — over 52,000 origin-destination pairs — feasible in seconds.

The timetable integrates directly: each scheduled class transition maps to an origin-destination pair. The graph topology makes bottlenecks immediately readable from the edge load distribution. The limitation of the approach is equally clear: within a corridor segment, there are no positions, no density waves, and no emergent local dynamics. For identifying building-scale load concentration, the approach is adequate; for analysing what happens physically at a specific doorway under high load, it is not.

Table 1 summarises the three approaches on the axes most relevant to this study.

Table 1: Comparison of modelling approaches on the axes relevant to this study.

Model	Comp. cost	Timetable integration	Spatial resolution
Macroscopic	Low	Poor	Corridor-level
SFM	High	Poor	Agent-level
Graph	Low	Natural	Segment-level

2.2 Simulated Annealing

The goal of the optimisation is to minimise corridor congestion by finding the room assignment that distributes student flows most evenly across the building’s edges. With n flexible rooms, the space of valid assignments has $n!$ configurations; at $n \approx 50$, this amounts to approximately 3×10^{64} possibilities. Exhaustive enumeration is impossible, and the objective landscape — edge load imbalance as a function of assignment — has many local minima, as pilot experiments with greedy assignment confirmed.

Simulated Annealing (SA), introduced by Kirkpatrick, Gelatt, and Vecchi [6], addresses this by accepting cost-increasing moves with a probability that decreases as the algorithm progresses. Starting from an initial assignment s_0 at temperature T_0 , each iteration proposes a neighbouring assignment s' obtained by swapping one to three room pairs. If the new cost $E(s') < E(s)$, the move is accepted unconditionally. If $\Delta E = E(s') - E(s) > 0$, the move is accepted with probability

$$P(\Delta E, T) = \exp\left(-\frac{\Delta E}{T}\right). \quad (3)$$

This Metropolis criterion allows the algorithm to escape local minima during the high-temperature phase. As the temperature decreases geometrically, $T(t) = T_0 \cdot \alpha^t$, the acceptance threshold tightens and the search converges. Parameters used here are $T_0 = 100$ and $\alpha = 0.98$; the algorithm reaches quasi-deterministic behaviour after approximately 500 iterations.

The cost function penalises both overall inequality in corridor load and extreme local bottlenecks:

$$E(s) = \sigma(\{n_e\}_{e \in E}) + \lambda \cdot \max_e (n_e - \bar{n} - 2\sigma)_+, \quad (4)$$

where n_e is the passage count on edge e under assignment s , $\sigma(\cdot)$ is the standard deviation of edge loads, $\bar{n}$ is the mean, and $\lambda = 10$ weights the penalty for edges exceeding the mean by more than two standard deviations. A lower $E(s)$ corresponds to a more evenly distributed load.

Three alternatives were considered and rejected. Integer Linear Programming would yield an exact solution but has exponential worst-case complexity for $n > 30$, ruling it out at the present problem size. Genetic algorithms incur higher computational cost without offering better solution quality at this scale. Greedy heuristics are fast but consistently became trapped in local minima in pilot experiments conducted before SA was adopted. SA offers the best trade-off between solution quality and running time for $n \approx 30$ flexible rooms.

3. Methods

3.1 Behavioural Experiment

Controlled experiments were conducted in the gymnasium of the current Mondial College Leuvenbroek building. A camera was mounted on a gymnastics bar at height $h \approx 7.5$ m, providing a near-orthographic overhead projection suitable for position tracking. Corridor and junction configurations — straight sections, a 90° corner, a door frame, and a T-junction — were constructed from gym cabinets and mats. Camera intrinsics (lens distortion) were calculated once per camera using a moving checkerboard calibration target and OpenCV; gymnasium floor markings served as the external scale reference.

Participants were $N = 208$ students aged 11–19, recruited via invitation to mentors. Participation was voluntary and data were collected under informed consent in accordance with AVG guidelines¹. Sessions ran in groups of 15–60 students, combining lower and upper school years.

Seven scenarios were recorded:

1. Straight corridor, unidirectional flow.
2. Straight corridor, bidirectional counterflow (two equal-sized groups).
3. Counterflow with a table-and-chair obstacle in the corridor.
4. Counterflow with a door construction in the corridor.
5. Counterflow through a 90° corner.
6. Corner with a stationary person at the apex.
7. T-junction with three groups converging simultaneously.

Scenarios 1–7 were also run with an instruction to use a smartphone during transit as an overlay condition, allowing the attentional friction effect to be isolated across configurations.

Fully automated trajectory tracking was attempted but abandoned. Group dynamics produced microscopic-level noise — occlusion, overlapping bounding boxes during cluster passage — that prevented reliable automated extraction of individual trajectories. Instead, video footage was analysed manually to develop a behavioural taxonomy for macroscopic calibration.

Free-walking speed was measured separately. Forty-one students walked individually along an unobstructed straight path at their own pace, with no other students present. Position over time was extracted from the calibrated video frames using the same tracking approach. The resulting $N = 41$ speed measurements provide the baseline parameter v_0 used in all downstream models.

3.2 Building Graph Construction

Architectural BIM drawings of the new building were obtained from the facility manager, covering all floors and including room function labels. A custom Python application was developed to convert these drawings into a simulation-ready network. The user annotates corridors, junctions, staircase landings, and room entrances interactively on the imported floor plan; the software automatically assigns each edge a traversal weight based on measured length and the empirically established free-walking speed. Staircase edges receive an additional cost correction reflecting reduced speed on stairs. Corridor width was recorded as an attribute for capacity reference during results interpretation; it was not used as a limiting constraint

¹the Dutch implementation of the GDPR

in the graph simulation itself. Width-based capacity limits are planned as an active constraint in the upcoming microscopic simulation.

Vertices V represent discrete locations: 51 classrooms, 74 secondary endpoints (storage rooms and dead ends), 113 corridor segment nodes, and 150 junctions where three or more paths meet, giving $|V| = 388$. Edges E encode passable connections between adjacent vertices. The final graph has $|E| = 427$ weighted edges. Corridor-segment nodes with exactly two connections — included for geometric accuracy but redundant computationally — were contracted in the simulation graph while preserving all distance information.

The building has two floors. Each floor was encoded as a separate sub-graph; staircase nodes connect the two layers. Of the building's seven staircases, five are accessible to all students.

A site visit to the new building before occupation confirmed observations already visible in the BIM maps. The central staircase exits directly into the main hall and is prominently located adjacent to the auditorium; it is the most natural inter-floor route for the majority of students. It has a funnel geometry: the base is wide enough for approximately five persons abreast, but the upper landing narrows to roughly 0.8 m, accommodating two at most. The other staircases are narrower than those at the current Leuvenbroek location. Besides that, the new building offers five publicly accessible staircases (of seven total) versus the effective 16 of the current four-floor, four-shaft building.

3.3 Timetable-Driven Simulation

The full weekly school timetable used for simulation is the current Leuvenbroek schedule, translated to the new building's room layout. It was extracted from the scheduling system for a representative regular week (no excursions, no abbreviated days; routine lesson cancellations were included). Student identification numbers were hashed before processing in accordance with AVG requirements. Parsing the timetable produced more than 52,000 origin-destination pairs per week, each specifying which room a student must travel from and to at a given lesson change.

For each origin-destination pair, Dijkstra's shortest-path algorithm was applied to the building graph. Repeated pairs — the same room combination recurring across the week — were cached to avoid redundant computation. The output is a per-edge passage count aggregated across all lesson changes in the full school week, giving a corridor load distribution from which bottleneck edges can be identified directly.

3.4 Room Allocation Optimisation

The state space for optimisation is the set of bijective mappings $\phi : L_{\text{old}} \rightarrow L_{\text{new}}$ from old room labels to new physical rooms. Of the 51 rooms in the new building, 21 are fixed by function — gymnasias, science laboratories, music rooms, and other specialist spaces cannot be freely reassigned — leaving $n_{\text{flex}} = 30$ flexible rooms. One old room label must map to the null room ($\rightarrow \emptyset$) to account for a capacity difference between buildings; the algorithm determines which assignment makes this optimal.

Each candidate assignment s is evaluated by running the full timetable simulation on the resulting graph and computing the cost function defined in Equation 4 with $\lambda = 10$. Five metrics are recorded per evaluation: mean edge load, standard

deviation of edge load, maximum edge load, bottleneck count (edges with load $> \bar{n} + 2\sigma$), and the Gini coefficient of the load distribution.

A neighbour assignment s' is generated by swapping 1–3 randomly selected room pairs. The SA temperature schedule uses $T_0 = 100$ and $\alpha = 0.98$, running for a maximum of 500 iterations. Optimisation is initialised with a greedy warm start: 20 random assignments are evaluated and the best serves as s_0 , saving approximately 50–100 iterations relative to a fully random start. One cost-function evaluation takes approximately 18 seconds on a standard laptop (Intel i5, Python); a full 500-iteration run takes approximately 2.5 hours. Convergence was validated by running five independent instances with different random seeds; best scores across runs differed by less than 3%.

Hard constraints were handled outside the optimiser: the facility manager fixed those assignments manually, and SA then searched over the 30 remaining flexible rooms.

3.5 Microscopic Simulation Framework

The graph model identifies where load accumulates at building scale but cannot represent how congestion develops locally — queue formation at a doorway, density waves propagating back through a corridor, or the physical consequences of the funnel staircase geometry. Phase 3 addresses this using JuPedSim [7], a research-grade C++/Python simulator in which each student's motion is governed by the Social Force Model. Behavioural parameters derived from the Phase 1 experiment — effective free-flow speed and corner-clipping bias — are incorporated as calibration inputs. Group cohesion parameters are also being explored, though the extent to which JuPedSim can faithfully represent the friend-group dynamics observed in the experiment is a question the Phase 3 work will need to address empirically.

JuPedSim needed a different kind of building description than the Phase 2 graph did: actual walls, room outlines, doors, pillars and obstacles, rather than a list of weighted edges. A second piece of engineering, separate from the Phase 2 annotation tool, supplies this. Each floor stores one registry of named vertices, and every room, wall, door, pillar, obstacle, breakline and inter-floor transit zone is defined only as references into that registry, never as raw coordinates. Doors are anchored to a position along their parent wall rather than to a fixed point in space, so repositioning the wall does not invalidate the door cut into it. The structure checks its own consistency — every referenced vertex must exist, and every inter-floor connection must have a matching zone on the destination floor — and it serialises losslessly to JSON. This representation currently supplies the geometry for the JuPedSim runs and for the building-rendering tool being built alongside them; folding the Phase 2 graph into the same representation, so that one edit to the floor plan reaches both models automatically, is planned but not yet built.

JuPedSim does not support multi-floor buildings natively. The solution adopted here runs one simulation instance per floor in parallel. Custom stair-transition logic monitors agents approaching staircase nodes and transfers them between floor instances at the appropriate time, preserving momentum and timing. A cross-floor visibility graph handles global routing decisions so that agents can plan paths that span both floors before being injected. The real timetable then drives agent releases: at each lesson change, agents are injected at their source rooms and routed to their destinations according to the

schedule, with a randomised v_0 based on the normal distribution from Phase 1.

Phase 3 is ongoing. The contribution at the current stage is methodological: the multi-floor parallelisation architecture, the cross-floor routing graph, and the BIM-to-JuPedSim conversion pipeline are the current primary outputs. Results will be demonstrated live at EUCYS 2026.

4. Results

The results reported here reflect the graph-theoretic model (Phase 2) and the physical experiment (Phase 1). Phase 3 microscopic simulation results for the entire building are not yet available.

4.1 Behavioural Characterisation

The phenomena below come from manual frame-by-frame analysis of the Phase 1 footage (Section 3.1) rather than from adult-population baselines in the literature, and they supply the calibration parameters used in the network and microscopic models that follow.

Free-walking speed

Free-walking speed, measured from $N = 41$ individual trials on an unobstructed straight path, was $\mu = 1.47$ m/s ($\sigma = 0.19$ m/s). The distribution is approximately normal, with most observations falling between 1.3 and 1.7 m/s. The value is consistent with published measurements for pedestrians [1] and serves as the baseline v_0 for all model components.

Social clustering

Group cohesion was present and consistent across every multi-person scenario. Within a few seconds of entering the corridor, students sorted themselves into friend groups that moved as coherent units. Groups walked abreast, reducing the usable corridor width proportionally to group size. At the T-junction, groups followed the direction chosen by the first member without independent re-evaluation; this produced rapid, committed splits rather than the gradual flow distributions an independent-agent model would predict. Routing decisions were correlated within groups, concentrating load on the edges that the group leader selected.

Corner-clipping

Across all corner scenarios, students preferentially took the geometrically shortest path through the turn, cutting the inside of the bend. The inside of the corner accumulated the highest local density; the outside remained consistently underused. The pattern appeared without instruction and without apparent awareness by participants. The consequence for capacity estimation is that the effective usable width at a corner is lower than the nominal corridor width, and the density peak occurs toward the inside of the turn.

Attentional friction

Scenarios with phone use produced observably slower throughput and increased stop-and-go behaviour compared to the control condition. Students using phones exhibited lateral drift and made unexpected stops without apparent awareness of downstream disruption. The effect was most pronounced at the corner, where a distracted student pausing at the apex — amplified by the stationary confederate — produced a near-complete stall in the surrounding flow. In terms of calibration,

this is encoded as a reduction in effective free-flow speed and an increase in the headway that distracted agents maintain.

Door passage

Door passage was bursty and asymmetric. Friend groups passed as clusters: once the first member entered the doorway, the remainder followed without pausing or alternating with the opposing stream. Alternating passage — which would have been more efficient for total throughput — was never spontaneously adopted. Bidirectional door use produced brief but complete blockages until one stream yielded.

Waiting classes

Classes waiting outside locked classrooms caused corridor blockage disproportionate to their number. Waiting students act like an ideal gas; they expand to fill all available space — and blocked both entry to the classroom and passage through the adjacent corridor segment. Widening the corridor did not resolve this: wider corridors were simply filled more completely. The single effective remedy observed was timely opening of the classroom door. A corridor wing in the new building where room entrances on both sides are closely spaced will be particularly susceptible to this effect.

Survey and interview findings

A student survey confirmed that lesson changes and staircase congestion are the primary self-reported pain points. Students noted that unwritten norms such as keeping to the right are routinely ignored and proposed a range of interventions, including staggered lesson times, fewer internal doors, and wider corridors.

Structured interviews with students, teachers, and the timetable coordinator produced a consistent finding across all three groups: longer lessons of approximately 60 minutes (from 45 minutes) are the most feasible near-term intervention. Reducing the number of daily lesson changes directly reduces the number of peak congestion events without requiring physical changes to the building or major restructuring of the timetable. Teachers unanimously rejected fixed classrooms for students, citing concerns about territorial behaviour and reduced pedagogical flexibility; they expressed clear support for fixed classrooms for teachers, citing reduced workload. The timetable coordinator confirmed that large structural reorganisation of the timetable is not feasible given the scheduling software, teacher availability constraints, and legal requirements on lesson distribution across the week.

4.2 Geometric and Topological Analysis

The building graph ($|V| = 388$, $|E| = 427$) drove simulation of the full weekly timetable, producing more than 52,000 origin-destination pairs. The resulting edge load distribution for the original room assignment is summarised in Table 2.

The distribution is highly unequal: the most-loaded edge carries 6,513 passages per day, approximately 6.6 times the mean. The Gini coefficient of 0.585 indicates substantial concentration of flow on a small fraction of edges.

The dominant bottleneck is the central hall. Across all simulated lesson changes, 61.99% of all student paths pass through this space. The cause is topological: the building has two classroom wings connected via a single central hub. Any student moving between floors, between wings, or from either wing to the entrance must traverse the hall. No room

Table 2: Edge load statistics for the original room assignment (values per school day).

Metric	Value
Mean edge load $\bar{n}$	986 passages
Std dev σ	1,262 passages
Maximum edge load	6,513 passages
Minimum edge load	13 passages
Bottleneck edges ($> \bar{n} + 2\sigma$)	18
Gini coefficient	0.585

reassignment can alter this; the bottleneck is a property of the graph structure, not of which room is assigned to which number.

The first floor carries substantially lower load than the ground floor. It is accessible only via staircases, and the distribution of traffic across the five available staircases is uneven, as the central staircase — the most prominently positioned, opening directly into the main hall — absorbs a disproportionate share of inter-floor movement.

Comparing staircase provision between buildings requires accounting for floor count. The current Leuvenbroek building has four floors and four staircase shafts, providing effectively 16 staircase flight segments in operation. The new building has two floors and five staircases, providing 5 flight segments — a net reduction in capacity despite the higher staircase count. Staircase widths in the new building are also narrower. Combined with the uneven load distribution, this means inter-floor passage will be meaningfully more constrained than at the current location.

The funnel staircase compounds this picture. Clearly visible in the BIM maps and confirmed on-site, it exits directly and prominently into the hall — the most loaded space in the building — making it the primary inter-floor route for the largest share of students. Its base is wide enough for approximately five persons abreast, but the upper landing narrows to roughly 0.8 m, concentrating the densest flows at the narrowest point.

4.3 Room Allocation Optimisation

The original room assignment proposed by the facility manager received a cost-function score of $E(s_{\text{orig}}) = 132,675$. The Simulated Annealing optimiser improved this to $E(s_{\text{opt}}) = 119,144$, a reduction of 10.2%. The algorithm stabilised after approximately 400 iterations; no further improvement was observed in the remaining 100.

Table 3 compares the five load metrics between the original and optimised assignments.

Table 3: Comparison of edge load metrics of an entire week between original and SA-optimised room assignments.

Metric	Original	Optimised	Change
Mean edge load	4,880.96	4,868.07	−0.3%
Std dev of edge load	6,145.52	5,864.99	−4.6%
Maximum edge load	29,825	27,926	−6.4%
Bottleneck edges	18	14	−4
Gini coefficient	0.5848	0.5675	−3.0%

The mean edge load is unchanged to within 0.3%, as expected: the total number of student movements is fixed regardless of

room assignment. The optimisation redistributes load rather than reducing it in aggregate. The standard deviation falls by 4.6%, the maximum by 6.4%, and the bottleneck edge count from 18 to 14. The Gini coefficient decreases from 0.5848 to 0.5675, a 3.0% reduction in load inequality.

Five independent SA runs with different random seeds converged to cost scores within 3% of one another, but to different room configurations. There is no unique optimal assignment; multiple configurations achieve comparable scores. The pattern common to the best configurations is that rooms assigned to the same subject are not placed in adjacent locations. The explanation is that a student who has maths in one room rarely has maths again in the immediately adjacent room for the next lesson; placing subject rooms apart slightly reduces average travel distance.

The practical search space is further constrained by hard requirements: ventilation, acoustic isolation, teacher co-location, and pedagogical clustering reduce the number of truly free assignments to fewer than the nominal 30. A two-stage approach is indicated — architects fix hard constraints first, then SA optimises over the remaining free rooms.

Three factors explain why larger improvements are not achievable. First, the building’s topology forces the majority of routes through the central hall regardless of which room has which label. Second, the existing timetable is already approximately balanced: each room sees a similar number of occupants across the week, leaving no systematic pattern that a reassignment could exploit. Third, only 30 of 51 rooms are freely assignable. A 3% Gini reduction is a genuine improvement, but room allocation is a complementary intervention, not a solution to topologically determined congestion.

5. Conclusions

The framework developed here — empirical behavioural characterisation, graph-theoretic timetable-driven simulation, and SA-based room optimisation — is practically viable for pre-occupancy crowd analysis and produces findings that are directly actionable before the building opens. The central result is that pedestrian congestion in this school is primarily determined by spatial topology, not by room assignment or timetable structure. The building’s two-wing layout, connected through a single central hall, concentrates 62% of all student paths through one space regardless of any organisational decision. No reallocation of rooms can change this.

Model choice. Among the three modelling classes evaluated, the graph-theoretic approach is optimal for building-scale, timetable-driven crowd analysis. It runs in $O(N \log |V|)$ time, integrates directly with timetable data, and makes bottlenecks readable from edge load counts. The SFM was implemented and tested; at $N \approx 1,200$ it was computationally infeasible as a whole-building model at the current phase, though ongoing improvements in optimised implementations and computing resources are making this feasible for Phase 3. Macroscopic fluid models yield useful theoretical capacity bounds but fail as routing models in corridors 2–4 m wide with heterogeneous destinations.

Geometric analysis. The baseline load distribution has a Gini coefficient of 0.585 and a peak-to-mean ratio of 6.6. The

central hall has a passage probability of 61.99%, imposed by the building's bone-shaped topology. Staircases in the new building are narrower and less evenly loaded than at the current Leuvenbroek location; inter-floor capacity is effectively reduced despite a higher staircase count. Keeping every staircase accessible matters more here than it would in a more evenly loaded building: the outer staircases are the only relief for a hub that is already carrying far more than its share. The move is months away; this is still actionable.

Behavioural characterisation. The three behavioural friction factors — social clustering, attentional friction, and corner-clipping — were confirmed across all relevant scenarios. All three enter the model as calibration parameters. The most consequential finding not captured by any geometric or topological analysis is the waiting-class effect: students blocking corridors outside locked classrooms do so completely and independently of corridor width. Timely classroom opening is a more effective remedy than corridor widening. The free-walking speed was $v_0 = 1.47 \text{ m/s}$ ($\sigma = 0.19 \text{ m/s}$), consistent with adolescent literature values.

SA optimisation. SA reduced the cost-function score by 10.2% and improved load-distribution metrics by 3–6% relative to the original assignment: Gini from 0.585 to 0.568, maximum edge load by 6.4%, bottleneck count from 18 to 14. No unique optimum exists; five independent runs converge to similar scores but different configurations. Room allocation is a complement to other interventions, not a substitute for them. The three-factor ceiling — topology, timetable uniformity, and restricted search space — limits further improvement.

Effective measures. Three categories of intervention are supported by the combined findings, ordered by expected impact.

1. *Architectural.* Keep the outer staircases accessible: closing them would funnel their share of inter-floor traffic onto an already overloaded central hub (Section 6.2). Address the funnel staircase: the narrowing from five-person width to approximately 0.8 m at the upper landing is a design problem at the building's most trafficked inter-floor connection. Widen or relocate the landing if structurally feasible.
2. *Organisational.* Introduce longer lesson blocks of approximately 60 minutes — a measure endorsed by all three interviewed stakeholder groups and the only intervention that directly reduces the number of peak congestion events per day. Ensure classrooms are opened before lesson-change bells ring; the waiting-class effect is large, controllable, and currently unaddressed.
3. *Behavioural nudging.* Floor markings and directional signage directing students toward the first floor and the outer staircase routes redistribute load from the most congested segments. A new building, before habitual routes have formed, is the optimal moment to establish these patterns.

Post-occupancy validation — planned for September 2026 when the building opens — will compare simulated load maps against on-site pedestrian counts at critical locations. The

methodology is building-agnostic: any facility with a floor plan and a timetable can be analysed with the same pipeline before architectural decisions become fixed.

6. Discussion

6.1 Limitations

Four limitations of the present study warrant explicit statement.

Deterministic routing. Every student is routed via the Dijkstra shortest path. Real pedestrian behaviour deviates from shortest-path predictions for social reasons — students may take a longer route to walk with a friend, avoid a certain staircase, or retrace a familiar path. The graph model does not represent any of these preferences. The partial mitigation is that the study concerns a building nobody has yet occupied: habitual routes and spatial preferences require repeated experience to form, so shortest-path routing is a reasonable approximation for the initial weeks of occupation, when the predictions will matter most.

Static timetable. The timetable is treated as a fixed input to the optimisation. Jointly optimising room assignment and timetable structure — for example, staggering which subjects change rooms at the same time — would in principle reduce peak loads further. The joint problem is combinatorially intractable and organisationally constrained by scheduling software requirements, teacher availability, and legal prescriptions on how lessons must be distributed across the week. The timetable coordinator confirmed these constraints directly; treating the timetable as fixed is not merely a modelling simplification but a reflection of what is actually variable.

No within-corridor dynamics. A network-flow model assigns passage counts to edges; it cannot represent what happens along an edge — queue formation at a doorway, a density wave that propagates back from a bottleneck, or the physical effect of the funnel staircase narrowing. This is the structural limitation of the graph approach, not a correctable oversight. Phase 3 exists precisely to address it: the JuPedSim microscopic framework, calibrated with the Phase 1 parameters, will produce per-location agent densities and velocity fields that the graph model cannot.

Behavioural parameters from the current building. The calibration parameters come from a constructed scenario in the current Leuvenbroek gymnasium, not from real lesson changes in the new building, which does not yet exist to film. Some of what was measured may be an artefact of that setting rather than a feature of how students normally move: participants knew they were taking part in an experiment, the obstacles were assembled from gym equipment rather than being an actual corridor, and there was no real lesson to rush to. We cannot fully separate genuine behavioural tendencies from effects of the experimental setup itself, and treat the resulting parameters as reasonable estimates to be checked against the September 2026 post-occupancy data rather than as exact measurements.

6.2 Practical Implementability

The formal optimisation covers 30 of the building's 51 rooms. Within those 30, hard constraints further reduce the effective

search space: acoustic separation requirements prevent music and drama rooms from being placed adjacent to examination classrooms; ventilation infrastructure fixes science-laboratory positions; teacher co-location preferences cluster certain subject rooms. The SA algorithm was applied after these constraints had been fixed by hand, meaning the optimised configurations already satisfy the hard requirements. The 3–6% metric improvements reported in Section 4.3 are achievable without violating any practical constraint.

A less obvious trade-off emerged during the study. The optimiser’s preferred configurations tend to disperse subject rooms — placing rooms assigned to the same subject at different locations in the building, which reduces the chance that consecutive lessons require travel between adjacent rooms. The same dispersal, however, increases teacher travel distances between lessons. Longer teacher transit times delay classroom opening, and the behavioural results identify waiting classes outside locked rooms as the single largest avoidable source of corridor blockage — one that neither geometry nor topology predicts. A room configuration that marginally improves the load distribution on corridors may simultaneously worsen the waiting-class effect. There is no algorithmic resolution to this trade-off; it requires case-by-case judgement about which rooms to disperse and which to keep clustered.

The measure with the clearest path from finding to implementation is longer lesson blocks. All three stakeholder groups interviewed — students, teachers, and the timetable coordinator — independently identified lessons of approximately 60 minutes (as opposed to the current length of 45 minutes) as the most realistic near-term intervention. Reducing the number of daily lesson changes from, for example, eight to six directly reduces the number of peak congestion events per day by 25% without requiring any physical modification, software change, or budget allocation. The convergence of stakeholder views across groups with otherwise different priorities is itself a finding: the measure is likely to be accepted if proposed.

A short example shows the kind of decision this approach is useful for. Suppose the administration considered locking the two outer staircases during school hours, on the assumption that the central staircases make them redundant. Removing those edges and re-running the timetable simulation would show how much of that displaced traffic the already-overloaded central hub would have to absorb — a question this kind of model can answer quantitatively before a decision is made, rather than after.

6.3 Broader Implications

Pre-occupancy crowd flow studies in secondary schools are rare in the published literature. Most crowd dynamics research addresses transport hubs, stadiums, and public events — settings where routing is demand-driven, agents are largely anonymous, and peak loads are episodic rather than daily and predictable [4, 2]. The timetable-driven, socially structured, and density-periodic character of a school population is qualitatively different. The three behavioural friction factors quantified here — social clustering, attentional friction from smartphone use, and corner-clipping — have not been systematically characterised and calibrated for this population in the prior literature. They are not universal corrections that can be borrowed from research on adult commuters; their magnitudes,

and the specific ways they interact with corridor geometry, are properties of student peer-group dynamics that required direct measurement.

The methodology is not specific to Mondial College. The geometry framework — a BIM/CAD-to-graph pipeline with a single representation feeding both network and microscopic simulations — can be applied to any building for which architectural drawings and a timetable exist. Hospitals, university campuses, and large office buildings all share the property of having scheduled occupancy and deterministic routing, making the same modelling approach applicable.

Post-occupancy validation is planned for September 2026, when the building opens. Pre-occupancy models that are subsequently validated against ground truth are rare; the comparison will both test the model’s predictive accuracy and identify where the graph model’s structural limitations become consequential.

6.4 Future Directions

Phase 3 completion. The JuPedSim microscopic simulation framework is the immediate next step. At the building scale, with $N \approx 1,200$ agents and the full floor plan as geometry, it will produce findings that the graph model cannot: the rate at which queues form and dissipate at the funnel staircase upper landing, the extent to which density waves in the central hall propagate back into the corridor wings, and whether the doorway conflicts observed in the behavioural experiment scale in ways the per-edge passage count does not predict. These are not refinements to a known answer — they may reveal that the graph model’s load estimates are systematically understated at specific locations where local dynamics, not route choice, are the binding constraint. Optimised implementations of JuPedSim and improved computing resources make whole-building microscopic simulation feasible for this phase.

Multi-floor parallelisation as a reusable contribution. The solution developed here for JuPedSim’s lack of native multi-floor support — parallel per-floor simulation instances, custom stair-transition logic, and a cross-floor visibility graph for global routing — addresses a limitation that affects any multi-storey application of the tool. The implementation is not tied to this building’s specific geometry. Others working with JuPedSim in hospital buildings, university libraries, or transit stations face the same constraint; the architectural approach developed here offers a concrete starting point.

Post-occupancy data collection and model refinement.

Three measurement methods are planned at different cost and intrusiveness levels. Anonymised overhead camera analysis at the central hall, the funnel staircase, and two peripheral staircases will provide crowd density and flow rate as continuous time series across a school day, enabling direct comparison with simulated edge loads. WiFi-based aggregate position sensing — using signal strength from access points already present in the building — offers a lower-intrusiveness alternative that can cover a larger fraction of the building simultaneously. Direct manual observation at a small number of locations provides the ground-truth baseline against which both other methods are calibrated. Discrepancies between prediction and observation will drive iterative model refinement: parameters that are systematically wrong in one type of location identify either

a missing behavioural factor or a structural limitation of the graph model.

Digital twin as a long-term direction. Real-time crowd sensing coupled to a continuously running simulation would make it possible to detect forming queues before they become blocking events and to redirect students via dynamic digital signage — for example, indicating when the central staircase is at capacity and the outer routes are clear. Comparable systems operate in large airports and rail stations, where the sensor infrastructure and the operational benefit justify the investment. A school-scale implementation faces different constraints: sensors must be unobtrusive and AVG-compliant, the system must be interpretable by facility managers without data science training, and the latency between sensing and response must be short enough to be useful during a three-minute lesson-change window. This is a long-term research direction rather than a near-term plan, but the current study provides the simulation backbone on which such a system would rest.

Declaration of AI Assistance

The research design, data collection, mathematical analysis, and original manuscript were entirely conceived and executed by the authors. The text was drafted by the authors, with a Large Language Model utilized strictly as a supportive tool for English language editing, vocabulary refinement, and LaTeX formatting assistance during the translation process.

References

- [1] Ernst Bosina and Mark Meeder. *Simulation of Complex Pedestrian Facilities: Case Study of Lausanne Main Station*. NSL – Netzwerk Stadt und Landschaft. June 16, 2017. URL: <https://www.nsl.ethz.ch/en/simulation-of-complex-pedestrian-facilities-case-study-of-lausanne-main-station/> (visited on 11/30/2025).
- [2] Mehdi Moussaïd et al. “The Walking Behaviour of Pedestrian Social Groups and Its Impact on Crowd Dynamics”. In: *PLOS ONE* 5.4 (Apr. 7, 2010), e10047. ISSN: 1932-6203. DOI: [10.1371/journal.pone.0010047](https://doi.org/10.1371/journal.pone.0010047). URL: <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0010047> (visited on 11/07/2025).
- [3] Armin Seyfried, Bernhard Steffen, and Thomas Lippert. “Basics of Modelling the Pedestrian Flow”. In: *Physica A: Statistical Mechanics and its Applications* 368.1 (Aug. 2006), pp. 232–238. ISSN: 03784371. DOI: [10.1016/j.physa.2005.11.052](https://doi.org/10.1016/j.physa.2005.11.052). arXiv: [physics/0506189](https://arxiv.org/abs/physics/0506189). URL: <http://arxiv.org/abs/physics/0506189> (visited on 09/05/2025).
- [4] Dirk Helbing and Peter Molnar. “Social Force Model for Pedestrian Dynamics”. In: *Physical Review E* 51.5 (May 1, 1995), pp. 4282–4286. ISSN: 1063-651X, 1095-3787. DOI: [10.1103/PhysRevE.51.4282](https://doi.org/10.1103/PhysRevE.51.4282). arXiv: [cond-mat/9805244](https://arxiv.org/abs/cond-mat/9805244). URL: <http://arxiv.org/abs/cond-mat/9805244> (visited on 08/11/2025).
- [5] Edsger W Dijkstra, Laurent Beauguitte, and Marion Maisonnobe. “E.W. Dijkstra, 1959, A Note on Two Problems in Connexion with Graphs. Numerische Mathematik 1, p. 269271 Version Bilingue et Commentée”. In: ().
- [6] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. “Optimization by Simulated Annealing”. In: *Science* 220.4598 (May 13, 1983), pp. 671–680. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.220.4598.671](https://doi.org/10.1126/science.220.4598.671). URL: <https://www.science.org/doi/10.1126/science.220.4598.671> (visited on 12/02/2025).
- [7] Mohcine Chraïbi et al. *JuPedSim*. Version v1.4.2. Zenodo, May 21, 2026. DOI: [10.5281/ZENODO.1293771](https://doi.org/10.5281/ZENODO.1293771). URL: <https://zenodo.org/doi/10.5281/zenodo.1293771> (visited on 06/07/2026).